LA-UR-03-3466

Approved for public release;
distribution is unlimited.

Title: A Risk-Based Approach to Designing Effective Security Force

Author(s): Terry Bott
Steve Eisenhawer

Submitted to: 44th International Nuclear Materials Management Annual Meeting
July 13-17, 2003
Phoenix, AZ

A Risk-Based Approach to Designing Effective Security Force Training Exercises

T. F. Bott and S. W. Eisenhawer
Probabilistic Risk Analysis Group, D-11
Los Alamos National Laboratory

Abstract

The effectiveness of a security force in protecting a facility is often evaluated using training exercises that pit a group of simulated attackers against a security team. In the situation studied here, the testers assigned increasingly sophisticated facility knowledge to the attackers. Security managers while not wanting to ignore these attacks feared they would be forced to concentrate resources on unrealistic scenarios at the expense of more credible threats. The problem was investigated using the Logic Evolved Decision (LED) method. The results of the analysis demonstrated that, for highly protected facilities, there was a strong tradeoff between the adversary's risk of interdiction arising from the attempt to collect more information before the attack versus settling for less information and mounting an attack with a lower likelihood of success. Security risk is not necessarily minimized by focusing exclusively on highly enhanced attacks at the expense of more probable, but less enhanced attacks.

Introduction

The effectiveness of a security force in protecting a secure facility is often evaluated using training exercises that pit a group of simulated attackers against a security team. In the natural progression of such exercises, as the defenders become more effective, the testing body imposes more challenging security scenarios by asserting, for example, that the attackers can disable important security systems. This amounts to implicitly assigning potentially high levels of facility knowledge to the attackers. Not surprisingly, a security force's assessed effectiveness decreases under such conditions, and the defender is forced to substantially alter his capabilities and response in order to successfully resist the attacks. As the sequence of exercises progresses, the knowledge attributed to the attackers becomes highly unrealistic for most adversaries, with the attendant danger that security resources are concentrated on stopping unrealistic scenarios at the expense of more credible threats.

To address this issue, we used the Logic Evolved Decision (LED) approach[1] to study the tradeoff from an attacker's viewpoint between increasing his likelihood of success by carrying out enhancing actions, versus increasing his risk of interdiction by gathering the extra information and knowledge required to carry out the enhancing actions. In this paper we discuss the methodology utilized and present illustrative results.

Methodology

The probability of success for an attacker can be enhanced by certain types of knowledge. Examples of this knowledge includes details about the electrical power system for the facility security equipment, and information concerning the communications system used by the security forces. Using this knowledge, the attacker could carry out actions to degrade the security response to an attack. These actions are termed enhancements and lead to enhanced scenarios. The knowledge required to accomplish the enhancement is called enhancing information.

Collecting enhancing information requires additional effort on the part of the attacker and will tend to increase the attacker's risk of detection and interdiction prior to the attack. To assess this increased cost to the attacker we modeled the espionage processes an attacker could use to collect different types of enhancing information. These espionage scenarios were generated using deductive logic gate models called possibility trees.[2]

A possibility tree is a convenient graphical method for creating complex directed graphs through a systematic deductive process. The steps in the deductive process involve the use of conjunction and disjunction. A disjunctive deductive step expresses a parent entity as an exhaustive set of specific instances. A conjunctive deductive step decomposes a parent entity into a conjunction of disjoint entities that taken together produce the parent entity. The relationship between the parent entity and its children is expressed through a logic gate. In a possibility tree, gates include traditional Boolean logic gates as well as non-commutative gates expressing more complicated relationships such as causality.

A possibility tree for gathering electrical power information is shown in Fig. 1. The top node G1 is a statement of the objective of the attacker's actions: "*The attacker causes a security power outag*e." Below this node, the ways this objective could be accomplished are deduced in progressively greater detail with each new level in the tree structure. The type of logic gate predominantly used in this development is called a causal gate and represents the logic of a process. For example, the goal stated at node G2, "*The attacker causes a general power outage*" is accomplished by the steps: "*The attacker takes out the normal power supply*" and "*The attacker takes out the emergency power supply.*" The first of these steps, shown as a diamond symbol at G3, denotes a logical structure that is collapsed (not shown).

The action described at node G4: "*The attacker takes out emergency facility power,*" is a process accomplished by carrying out steps G5, "*The attacker obtains information about emergency power,*" and G6 "*The attacker carries out the attack on the emergency power system*." Under node G5, the possible sources of information (e.g., "from facility website") are enumerated. Under each source the measures designed to protect the information are enumerated. A logical description of an objective and how it can be carried out, including the preventive measures designed to defeat the information gathering is coded in the logic-gate tree thus constructed.

A well-defined manipulation of the information coded in this tree produces a path solution that is identical to the paths through the directed graph represented by the logic-gate tree. The path solution provides a list of scenarios for gathering electrical power information that the attacker can use to degrade the security system and enhance the attack. The paths also includes the preventives in place to inhibit the collection of the required knowledge. An illustrative path is:

The attacker causes a security power outage. The attacker causes a general power outage. The attacker takes out the emergency power supply. The attacker takes out emergency facility power. The attacker obtains information about emergency power from the facility website. The website protective measures include website access limited to facility personnel and limited electrical system detail on the website.

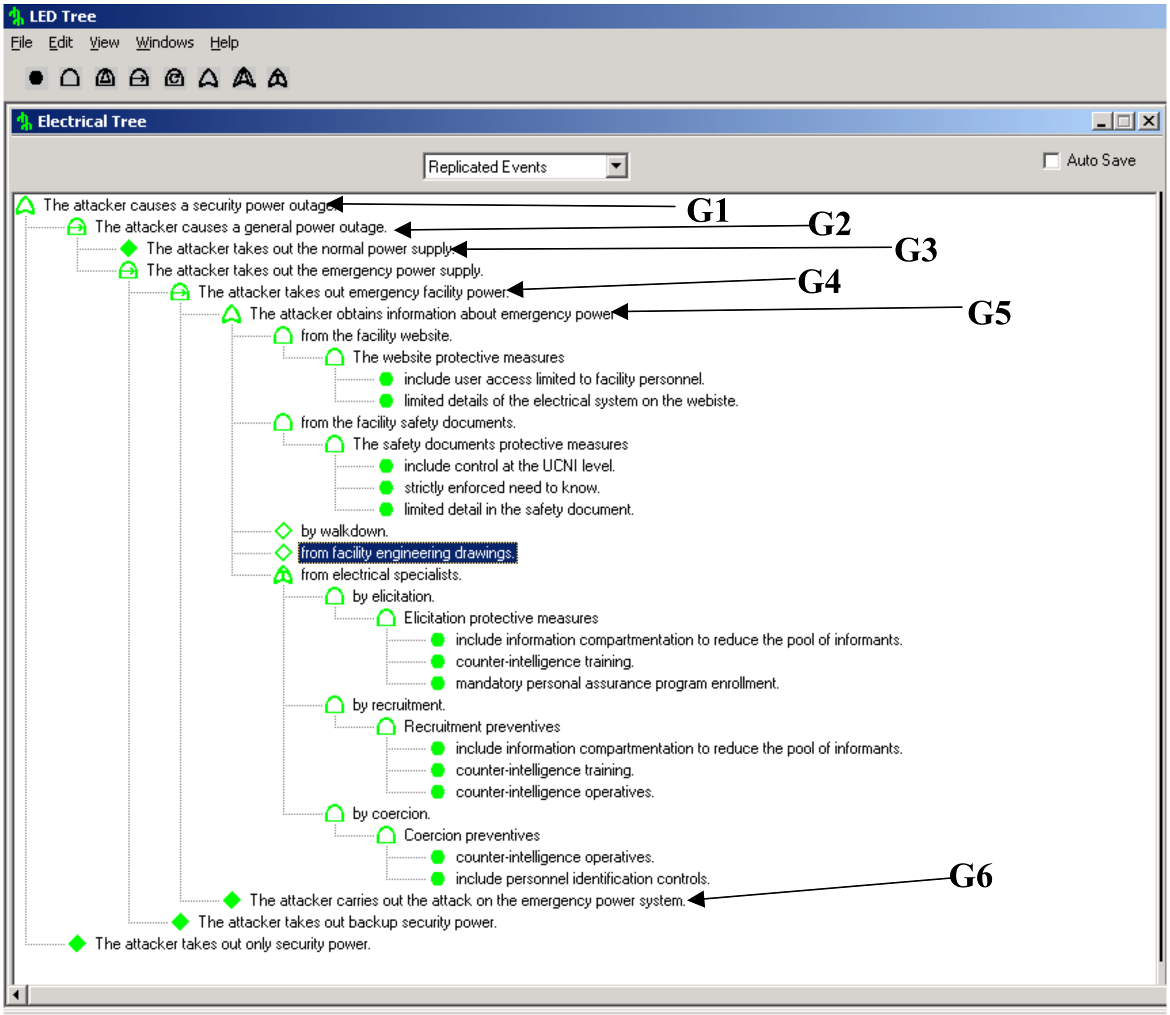

- The attacker causes a security power outage.
  - The attacker causes a general power outage.
    - The attacker takes out the normal power supply.
    - The attacker takes out the emergency power supply.
      - The attacker takes out emergency facility power.
        - The attacker obtains information about emergency power
          - from the facility website.
            - The website protective measures
              - include user access limited to facility personnel.
              - limited details of the electrical system on the webiste.
          - from the facility safety documents.
            - The safety documents protective measures
              - include control at the UCNI level.
              - strictly enforced need to know.
              - limited detail in the safety document.
          - by walkdown.
          - from facility engineering drawings.
          - from electrical specialists.
            - by elicitation.
              - Elicitation protective measures
                - include information compartmentation to reduce the pool of informants.
                - counter-intelligence training.
                - mandatory personal assurance program enrollment.
            - by recruitment.
              - Recruitment preventives
                - include information compartmentation to reduce the pool of informants.
                - counter-intelligence training.
                - counter-intelligence operatives.
            - by coercion.
              - Coercion preventives
                - counter-intelligence operatives.
                - include personnel identification controls.
        - The attacker carries out the attack on the emergency power system.
      - The attacker takes out backup security power.
  - The attacker takes out only security power.

Figure 1. Attack Enhancement Possibility Tree

The electrical system degradation portion of the scenario in black describes what the attacker does to enhance his likelihood of success. The information gathering process for the scenario is shown in text in red. The preventive measures associated with the information gathering process, shown in blue, provide a basis for evaluating the cost incurred by the attacker in gathering the enhancing information.

An important part of our approach was to determine the relative likelihood that different possible attack scenarios are attempted. We computed the attempt likelihood of attack scenarios by constructing a game. In this game, the players are the defender and the attacker. The object of the game is for the attacker to access the target and the defender to prevent this access. The game is played under incomplete information: the attacker does not know all of the details of the defenses, but can try to gain more details to enhance his probability of winning. Conversely, the defender does not know the extent of the attacker's knowledge, but tries to prevent access to helpful information and interdict the attacker during the information-gathering phase of the attack.

In this game we represented the probabilities and utilities associated with the game as linguistic variables and expressed uncertainty by treating the linguistic values of the variables as fuzzy sets. In our game, utility is measured by the linguistic variable **Attractiveness** and is a measure of the relative preference an attacker seeking maximum payoff would exhibit towards a set of attack scenarios.

Attractiveness is inferred based on a number of factors using an inferential model. The structure of the **Attractiveness** inferential model is shown in Fig. 2. Input variables are shown as dark-colored nodes. The light colored nodes are *If … then…* rule sets that show how values of their input variables combine to produce values of the node variable. The uncertainty is expressed using fuzzy set membership values, in contrast to a Bayesian network approach in which uncertainty in the values is expressed using conditional probabilities.

The inferential model considers the **Attractiveness** of each scenario from the attacker's viewpoint. **Attractiveness** is inferred from the variables "**Total Adversary Risk,**" "**Info Gathering Effort,**" and "**Final Success Likelihood.**" The most attractive scenarios for the attacker will be those that maximize the success likelihood while imposing acceptable costs and risks of interdiction during the information-gathering phase of the attack. The success likelihood includes both the success likelihood without any enhancements, and the effect of the enhancements.

To generate the **Attractiveness** estimate for an attack scenario that includes enhancements, values must be assigned for each of the input variables (the darker oblongs) in Fig. 2. These assignments are elicited from security and counter-intelligence experts. An example of a linguistic variable is "**Confidence in Electrical Info**" shown at the left top corner of the inference diagram as an input to "**Electrical Info Enhancement of Success Likelihood.**" The variable "**Confidence in Electrical Info**" can take on the linguistic values *low, moderate* or *high,* each of which is a fuzzy subset of the variable. The membership assigned to each of these possible fuzzy subsets reflects an expert's beliefs concerning the values for that scenario. The other input to the variable "**Electrical Info Enhancement of Success Likelihood**" is the variable "**Likelihood of Successful Elec Info Collection,**" which also takes on linguistic values *low, moderate* or *high*.

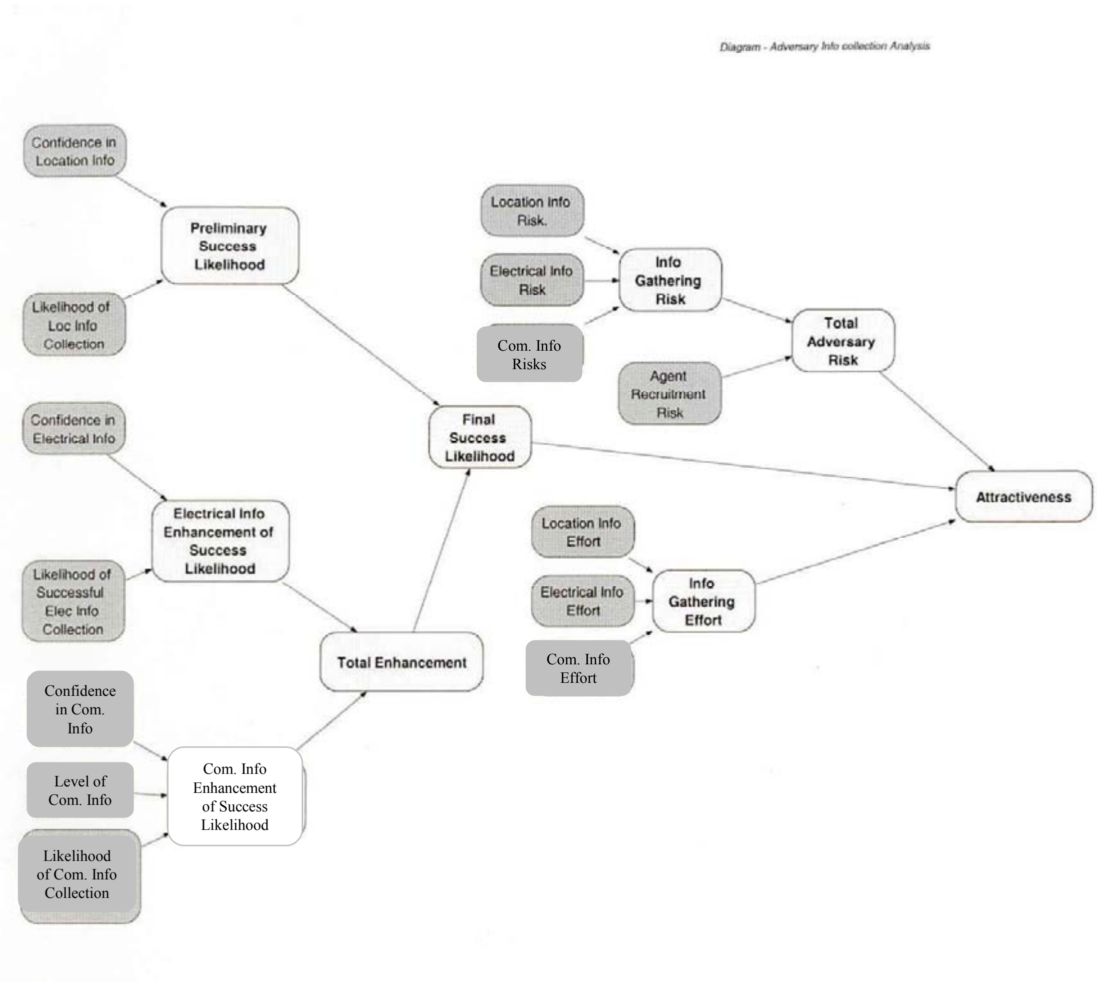
Diagram - Adversary Info collection Analysis
Confidence in Location Info
Preliminary Success Likelihood
Likelihood of Loc Info Collection
Location Info Risk.
Electrical Info Risk
Com. Info Risks
Info Gathering Risk
Total Adversary Risk
Agent Recruitment Risk
Confidence in Electrical Info
Final Success Likelihood
Attractiveness
Electrical Info Enhancement of Success Likelihood
Likelihood of Successful Elec Info Collection
Location Info Effort
Electrical Info Effort
Info Gathering Effort
Com. Info Effort
Total Enhancement
Confidence in Com. Info
Level of Com. Info
Com. Info Enhancement of Success Likelihood
Likelihood of Com. Info Collection

Figure 2. Inferential Model for Assessing Attack Attractiveness

Once values have been assigned to all the input variables in the inference diagram, the values propagate through the linked *If…Then…* rules to produce a value for the output of the inference model: "**Attractiveness**." An example of an *If…Then…* rule for "**Electrical Info Enhancement of Success Likelihood**" is shown in Fig.3. The analyst has great freedom in choosing the linguistic values and the *If…Then…* rules to capture the relationships between the factors in the inferential model. For example, in the rule shown in Fig. 3. High "Likelihood of Successful Elec Info Collection" and High "**Confidence in Electrical Info**" implies a value of *high* for "**Electrical Info Enhancement of Success Likelihood**" as outlined in yellow. Uncertainty in the results is propagated based on uncertainty in the inputs using fuzzy sets. The mathematics for this propagation is the Max-Min operation commonly used in fuzzy sets controllers.[3]

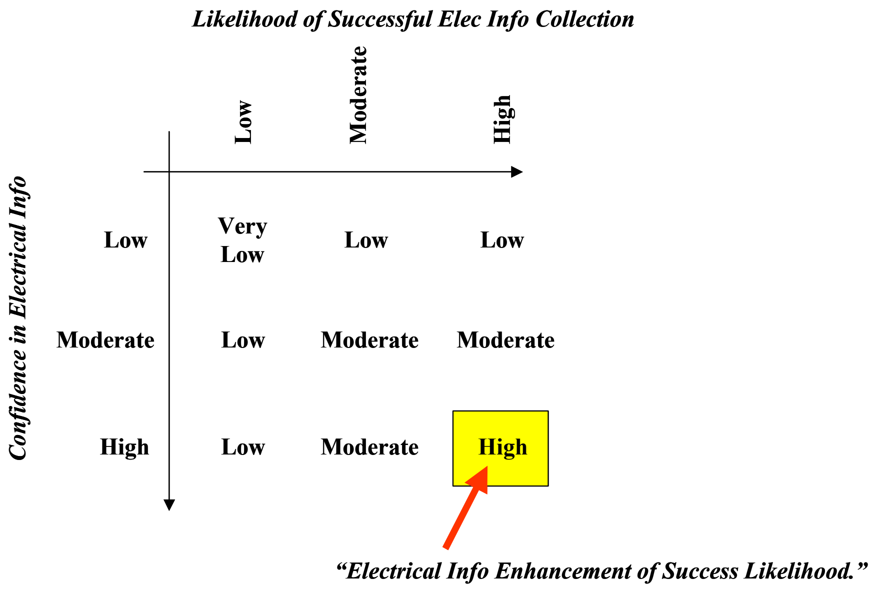
Likelihood of Successful Elec Info Collection
Confidence in Electrical Info

Confidence in Electrical Info \ Likelihood of Successful Elec Info Collection	Low	Moderate	High
Low	Very Low	Low	Low
Moderate	Low	Moderate	Moderate
High	Low	Moderate	High

"Electrical Info Enhancement of Success Likelihood."

Figure 3. Typical *If …* *Then …* Rule for Inferential Model

Illustrative Results

The specific details of the actual analysis are sensitive, but we can illustrate the kinds of analysis performed and the types of results obtained with a hypothetical example. In this example, an attacker wishes to access a protected material stored in a secure facility. This facility is protected by an armed guard force and elaborate physical security systems. These systems include protected electrical power sources for security equipment, secure communications that include broadcast and other elements, and a personnel security system that includes clearances, counter-intelligence programs, a system of security badges, compartmented information and a personal assurance program.

In actual studies, input data for our inferential model is collected from personnel knowledgeable about the electrical, communications and material tracking aspects of the facility, as well security, intelligence and counter-intelligence personnel. The experts we worked with found the linguistic inferences and the fuzzy set uncertainty representation to be a natural mode for expressing their ideas. In this example analysis we use representative values.

To carry out our example analysis we assume that all practical scenarios for gaining detailed information on disabling electrical power to security systems require that the adversary gain access to either documents or electrical specialists for the target facility. Controls on the documents require a compelling need in order to gain access. The electrical specialists are the most direct source of useful information for an attacker. Given access to the electrical specialists, the aiding information could be elicited, coerced or the specialists could be recruited. Elicitation carries high risks to the adversary, however, because the electrical specialists are very aware of the sensitivity of their knowledge, are very reluctant to share it with anyone and could become suspicious if anyone asks detailed questions.

Attractiveness values for the example are shown in Fig. 4. Each of the enhancement levels represents disabling of electrical (E) and communications (P) to some level. Three levels of communication disruption are: disruption of all communication channels (F), disruption of broadcast communications only (B) and none (N). For security-system electrical power disruptions, only the levels full (F) and none (N) are considered. Thus, E-F P-B implies the attacker gathered sufficient information to completely disable the security system electrical power s and disable broadcast communications capabilities. The axis labeled "measure" shows lists not only **Attractiveness**, but some interesting intermediate nodes in the inference diagram. The relative value is the centroid for each measure.[3] The bar graph shows that the most attractive scenarios (E-F P-N and E-N P-N) involve no or minimal extra information gathering on electrical or communications systems. Although in those cases there is less enhancement of the success likelihood, there is also a significantly reduced information gathering risk. In contrast, when full electrical and full or broadcast communication aiding information is gathered (E-F P-F and E-F P-B), there is a potential enhancement in the success likelihood from the attackers standpoint if the attack is carried out. But the attacker pays a heavy price before the attack in the form of an increased risk of discovery and interdiction during the information gathering effort.

The example analysis results are typical. The attack scenarios with the highest likelihood of success given an actual attempt also carried the highest risk to the adversary of prevention or interdiction before an attempt was even made. The interdiction and prevention likelihood were higher principally because of the actions required to collect highly detailed target and facility information. There was a strong tradeoff between the adversary's risk of interdiction arising from the attempt to collect more information before the attack versus settling for less information and mounting an attack with a lower likelihood of success.

Conclusions

When the only practical information gathering scenarios involve the attacker co-opting a select group of insiders who have access to facility details, the risk is greatly increased to the attacker. This is because the attacker does not have a large pool of potential collection targets and has a lower likelihood of locating a willing subject than when the potential pool is large. This risk is increased even more when information is effectively compartmentalized, because the attacker has to collect from members of several knowledge pools. The risk to the attacker is further increased when the members of the target knowledge pool are in a personnel assurance program, are aware of the importance of the information they hold, and are reticent about sharing it except to those with a compelling need to know. Compartmentalization of information and personnel assurance programs increase the chance that attempts to gain information will be recognized and resisted. The chance of counter intelligence operations that result in interdiction of the attacker are also increased significantly. These factors significantly increase the risk the attacker takes in trying to obtain more information to enhance the likelihood of success. This increased risk must be weighed by the attacker against the perceived enhancement achieved by gaining the extra information about the electrical power system.

These results suggest that realistic gaming best serves the objective of minimizing security risk in highly protected facilities. Some resources should be devoted to scenarios with significant attack enhancements that are less likely to be attempted, but more likely to succeed. But it is important to maintain readiness for more straightforward attacks, because those attacks are more likely to be encountered than attacks in which the adversary has full and detailed facility knowledge. The actual risk from the minimally enhanced attacks will often be greater than from highly enhanced attacks.

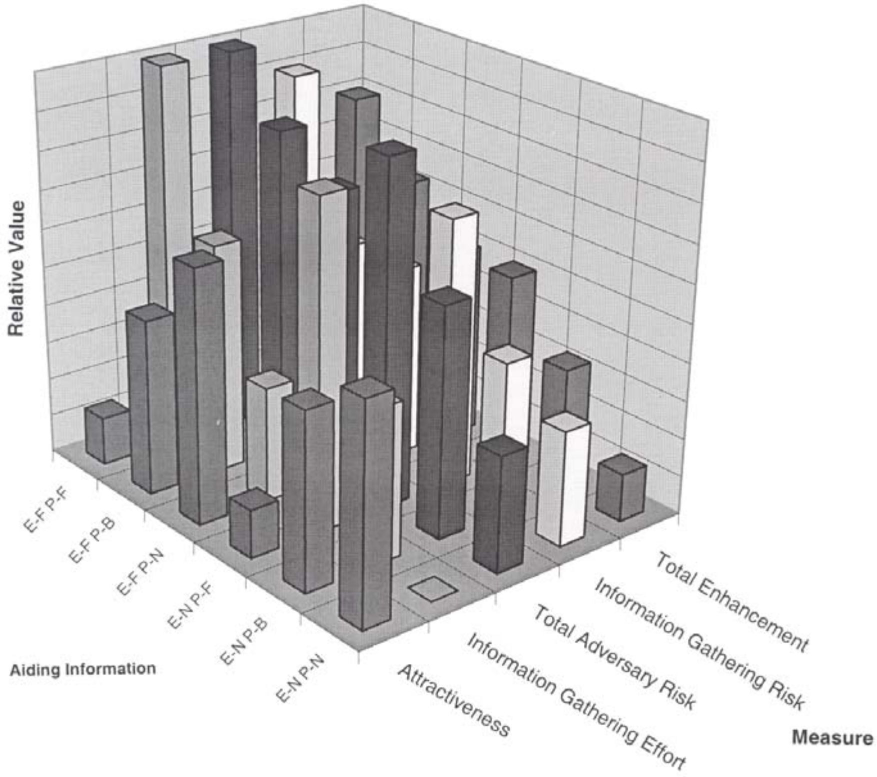

Figure 4. Typical Results for the Analysis

References

1. S. W. Eisenhawer and T. F. Bott, “Application of Approximate Reasoning to Safety Analysis,” *Proc. International System Safety Conference*, 1999 Aug., pp 374-382.
2. T. F. Bott, S. W. Eisenhawer, J. Kingson, and B. P. Key, “A New Graphical Tool for Building Logic-Gate Trees,” to appear, *ASME-PVP Annual Meeting*, Cleveland, Aug. 2003.
3. S. W. Eisenhawer, T. F. Bott, and R. E. Smith, “An Approximate Reasoning-Based Method for Screening High-Level-Waste Tanks for Flammable Gas,” **Nuclear Technology**, Vol. 130, 2000 June, pp 351–361.